

Application of State Analysis and Goal-Based Operations to a MER Mission Scenario

J. Richard Morris*

Michel D. Ingham[†], Andrew H. Mishkin[‡], Robert D. Rasmussen[§], and Thomas W. Starbird**
Jet Propulsion Laboratory, NASA, Pasadena, CA, 91109, USA

State Analysis is a model-based systems engineering methodology employing a rigorous discovery process which articulates operations concepts and operability needs as an integrated part of system design. The process produces requirements on system and software design in the form of explicit models which describe the system behavior in terms of state variables and the relationships among them¹. By applying State Analysis to an actual MER flight mission scenario, this study addresses the specific real world challenges of complex space operations and explores technologies that can be brought to bear on future missions. The paper first describes the tools currently used on a daily basis for MER operations planning and provides an in-depth description of the planning process, in the context of a Martian day's worth of rover engineering activities, resource modeling, flight rules, science observations, and more. It then describes how State Analysis allows for the specification of a corresponding goal-based sequence that accomplishes the same objectives, with several important additional benefits.

Nomenclature

<i>APXS</i>	=	Alpha Particle X-Ray Spectrometer
<i>ECam</i>	=	Engineering Camera
<i>Hazcam</i>	=	Hazard Avoidance Camera
<i>IDD</i>	=	Instrument Deployment Device
<i>MAPGEN</i>	=	MER Activity Plan Generator
<i>MB</i>	=	Moessbauer Spectrometer
<i>MDS</i>	=	Mission Data System
<i>MER</i>	=	Mars Exploration Rover(s)
<i>MI</i>	=	Microscopic Imager
<i>Mini-TES</i>	=	Miniature Thermal Emission Spectrometer
<i>Navcam</i>	=	Navigation Camera
<i>ODY</i>	=	Mars Odyssey Orbiter
<i>Pancam</i>	=	Panoramic Camera
<i>PMA</i>	=	Pancam Mast Assembly
<i>RAT</i>	=	Rock Abrasion Tool
<i>RML</i>	=	Rover Mark-up Language
<i>RoSE</i>	=	Rover Sequence Editor
<i>RSVP</i>	=	Rover Sequence Visualization Program
<i>Seqgen</i>	=	Sequence Generator
<i>sol</i>	=	Martian Solar Day

* Software and Systems Engineer, Planning and Sequencing Systems Group, M/S 301-250D.

[†] Senior Software Engineer, Flight Software Systems Engineering and Architectures Group, M/S 301-225, AIAA Member.

[‡] Principal Engineer, Planning and Execution Systems Section, M/S 301-250D.

[§] Principal Engineer, Systems and Software Division, M/S 301-225.

** Principal Engineer, Planning and Execution Systems Section, M/S 301-250D.

I. Introduction

As the complexity of space missions grows, operations risks are becoming a limiting factor. Mission requirements are no longer satisfied with such fault responses as “run to safe mode”. Coordination of multiple assets has itself become a common mission requirement, such as UHF communications relays between Mars landers and orbiters. And as the effort to realize NASA’s vision to return humans to the moon and on to Mars begins in earnest, human-robotic interaction will become a driving issue for the construction and maintenance of a lunar outpost. One of the most complex and ambitious space missions to date, the Mars Exploration Rovers (MER), two robotic explorers operating daily on the surface of Mars, has already taken the operations challenge to a new level.

The new challenge involves the simultaneous operability of multiple assets, the distributed operations required by concurrent missions and geographically dispersed expertise, and the scalability of current systems with respect to flight rules and operations personnel. A solution that addresses all of these challenges would incorporate a common architecture for distributed operations and model system state explicitly for operability². This paper presents State Analysis as such a solution.

The driving questions motivating this study include: How would a goal-based approach to operations benefit the MER uplink process? How would a Goal-based operations process work? Can State Analysis products provide additional benefit to the mission operations team?

The objective of this study is to initiate the incremental process by which we eventually answer those questions in concrete terms. Here we attempt to map a MER command sequence into a goal-based sequence (Goal Network), provide a comparison of the two operations approaches, and further refine our understanding of the scope, process, and tools associated with the Goal-based adaptation of a real system.

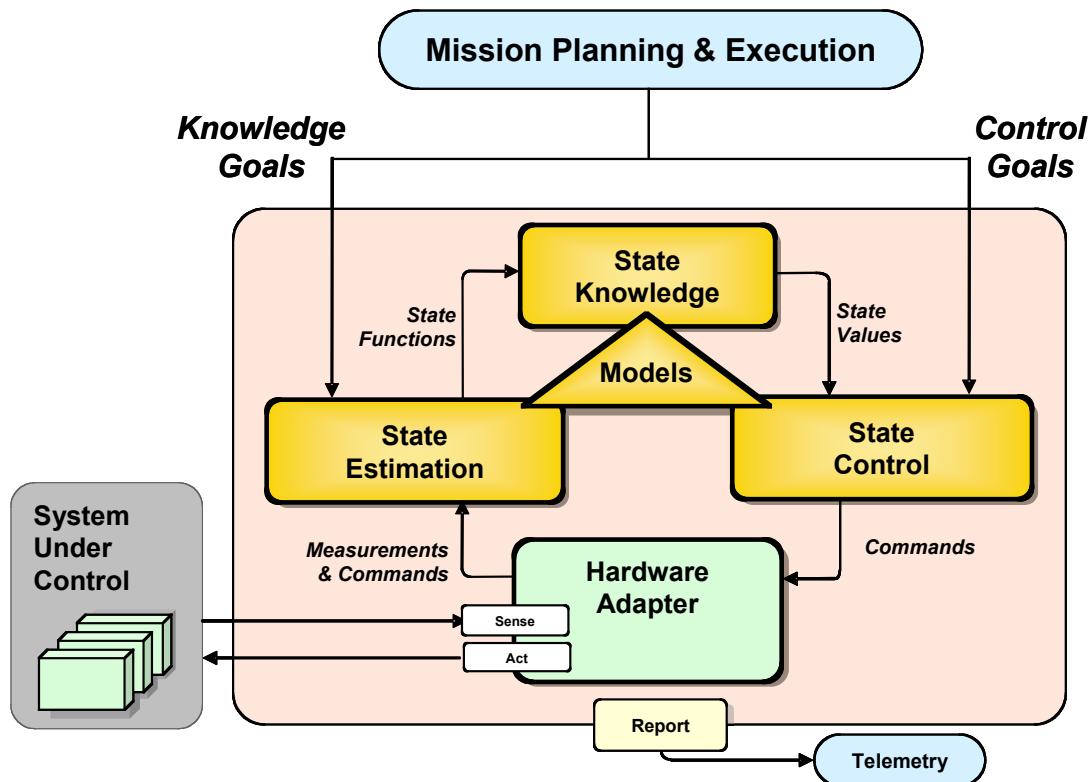


Figure 1. The state-based control architecture.

II. State Analysis

A. Introduction

State Analysis provides a process for capturing system and software requirements in the form of explicit models. State Analysis defines a uniform control architecture whose core principle is the explicit modeling of the state of the system under control. State Analysis goes on to provide a rigorous methodology by which the architectural elements in the control system (Fig. 1) are specified and documented. The resulting product set uniformly represents the systems engineers' understanding of the control system in terms of models of the system under control. The tight structural relationship maintained between these two domains provides a powerful architectural tool to both systems and software engineers, a key to managing operational complexity¹.

B. State-based Control Architecture

Several key features of the control architecture can be seen in Fig. 1. *Representation of State knowledge is explicit*, in the form of continuous-time State Functions, such that our best knowledge of the system's state is defined at any point in time. *State estimation is separate from state control*, which promotes objective assessment of system state and simplifies design through modularity. *Hardware adapters provide the sole interface between the system under control and the control system*, forming an explicit boundary of the control architecture. *Models are ubiquitous throughout the architecture*. Models are documented explicitly and inform the control system at each level of the architecture including state estimation, state control, and higher-level planning. *Control is performed through Goal-directed operation*, enabling the system as well as operators themselves to directly reason about control objectives. Goal-based sequences are derived directly from the model of the system under control, providing a verifiable link between requirements, design documents, and command loads. *The architecture provides for a straightforward mapping of the architectural elements into components in a modular software architecture*, such as the Mission Data System (MDS). Although a detailed description of MDS is beyond the scope of this paper, we should note that the application of State Analysis in conjunction with MDS serves to bridge the fundamental gap between the requirements on software specified by systems engineers and the implementation of these requirements by software engineers¹.

State Analysis can be further summarized by the following three main aspects:

- 1) State discovery and modeling
- 2) Control system design
- 3) Goal-based operations design

Aspects 1 and 3 are the focus of this paper.

C. State Analysis Methodology

State Analysis begins with the modeling of states of the system under control through an iterative process which incrementally constructs the model. The states include any physical value which we may want to control, which affect control, and/or which may need to be estimated. The states also include any data states we may want to manage. The steps in this process are as follows:

- 1) Identify needs – define the high-level objectives for controlling the system.
- 2) Identify state variables that capture what needs to be controlled to meet the objectives, and define their representation.
- 3) Define state models for the identified state variables – these may uncover additional state variables that affect the identified state variables.
- 4) Identify measurements whose values are affected by the identified state, and define their representation. Measurements are the input to the control system and exist at the interface to the system under control.
- 5) Define measurement models for the identified measurements – these may uncover additional state variables.
- 6) Identify commands that can influence the identified state, and define their representation. Commands are the output of the control system and exist at the interface to the system under control.
- 7) Define command models for the identified commands – these may uncover additional state variables.
- 8) Repeat steps 2-7 on all newly discovered state variables, until every state variable and effect we care about is accounted for.
- 9) Return to step 1, this time to identify supporting objectives suggested by affecting states (a process called 'goal elaboration,' described later), and proceed with additional iterations of the process until the scope of the mission has been covered.

These steps can be interwoven and performed concurrently at the discretion of the engineer.

D. State Effects Diagrams

In State Effects Diagrams, state variables are depicted as ovals, measurements as triangles, and commands as inverted triangles, as shown in Fig. 2.

State effects diagrams provide several important advantages. A single integrated model means consistency across project phases, tools, and roles. A member of the downlink team may be interested in assessing a trade between data volumes and data rates for communications, whereas a member of the uplink team may be more interested in trades between battery energy available and science observations. An engineer creating detailed elaboration and goal specifications during the formulation phase of a mission may be interested in state effects diagrams at a different level of detail than a member of the tactical operations team assembling plans on a daily basis.

State effects diagrams elevate design specification to a level that uplink personnel can understand in a timely manner. When implementing an activity a member of the uplink team will ask “what do I need to get this done?” This question is typically answered not by a tool, but by another operator. In effect, the team is going through the state discovery process every shift. This ad-hoc state discovery is error-prone and tedious. Why not capture the discovery process and shorten the planning process?

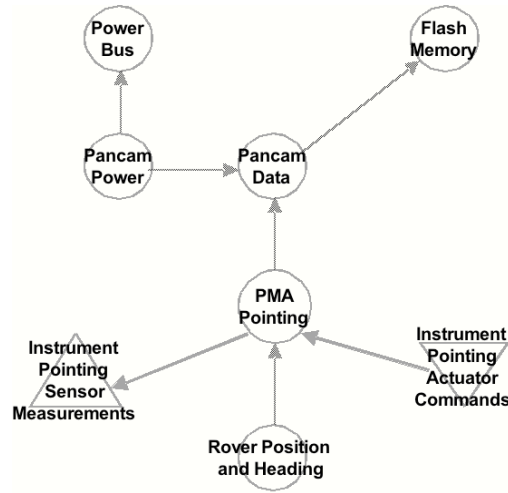


Figure 2. State Effects Diagram for Pancam-related State Variables.

E. Goal-based Operations

A Goal is an expression of operator intent. In State Analysis, this is expressed as a constraint on a State Variable's value history over a time interval. A Goal expresses *what* to do, not *how* to do it. The expression of a Goal network is more compact and inspectable than a traditional command sequence, making it easier to see interactions and conflicts between activities. This benefits operators and software alike, both of which are required to reason about directives and resolve conflicts between their probable and desired result. Because goals specify intent in terms of state, they are inherently closed-loop, giving both the operator and the software the ability to directly monitor the system for faults. Goals give the real-time control layer flexibility in achieving the goal.

The time interval associated with a goal is defined by a start and end *time point*, a potentially variable moment in time. Operators construct plans by instantiating goal types and interconnecting goal instances by either sharing time points or relating time points with temporal constraints, thus forming a *goal network*.

Because of the cause and effect relationship between states as modeled during the state discovery process, goals levied on one state variable may require supporting goals on other state variables which affect it. The process of expansion of a goal into supporting goals is directly informed by the state effects model and is called *elaboration*.

Goal elaborations are defined based on engineering judgment, our model of the system under control, and the following five rules:

- 1) A goal on a state variable may elaborate into concurrent control goals on directly affecting state variables (concurrent goals share the same start and end time points).
- 2) A control goal on a state variable elaborates to a concurrent knowledge goal on the same state variable (or they may be specified jointly in a single control and knowledge goal).
- 3) A knowledge goal on a state variable may elaborate to concurrent knowledge goals on its affecting and affected state variables.
- 4) Any goal may elaborate into preceding goals (typically on the same state variable). For example, a “maintenance”-type goal on a state variable may elaborate to a “transition”-type goal on the same state variable, which has an ending time point coincident with the starting time point of the “maintenance”-type goal.
- 5) A goal's elaboration can include uncertain temporal constraints to reserve time in the schedule for actions required by the goal.

Within these rules, engineers are free to exercise judgment on how exactly to achieve the desired ends. Flight rules and constraints, resource management, fault monitoring, and control are all accounted for within goals and goal

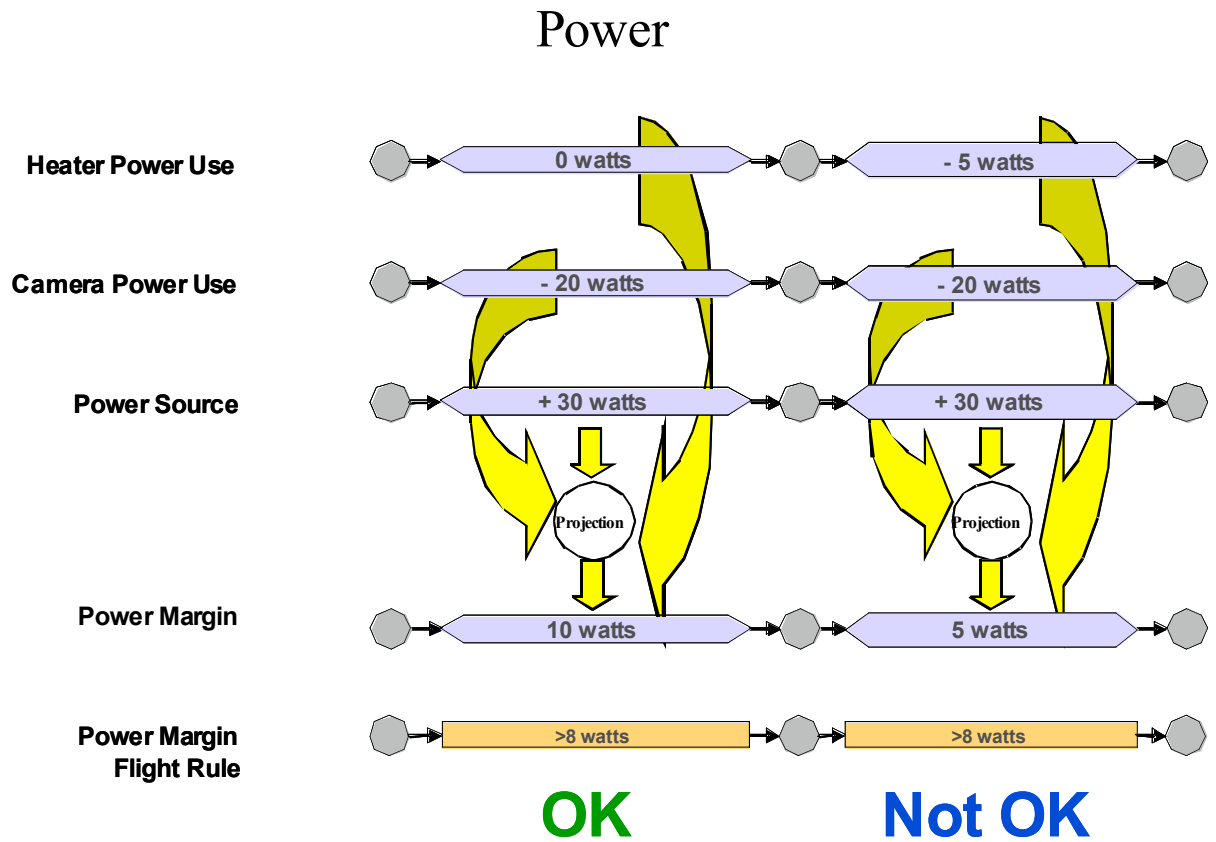


Figure 3. Projection for Power Goals.

elaborations. This consistent representation allows for direct reasoning and detection of conflicts between plans, and their likely result. Furthermore, given the model, state effects can be propagated not only across state variables, but also forward in time. Given a particular goal, and the control model, a controller can predict its likely performance over time. This process is called *state projection*. This gives the planning system the opportunity to discover conflicts before they occur. Since activities and resources are both expressed in terms of goals, projection becomes the process by which we model resource usage. By applying our elaboration rules, the effects of activities become goals on resources. This process can be automated and used onboard.

A simple projection example for power margin is shown in Fig. 3. In this example, the goal elaboration rules as applied to the state effects diagram result in goals on the state variables: heater power, camera power, and a power source. The model describing the effect of these state variables on the power margin state variable results in goals on the power margin state variable. These power margin goals can now be compared directly to flight rules, in the form of goals, to discover resource violations.

III. MER Operations Process

A. Overview

In 2003 NASA's Jet Propulsion Laboratory sent two robotic rovers to Mars. Surface operations for *Spirit* began on January 3, 2004 and for *Opportunity* three weeks later. Both rovers continue to return science data on a daily basis.

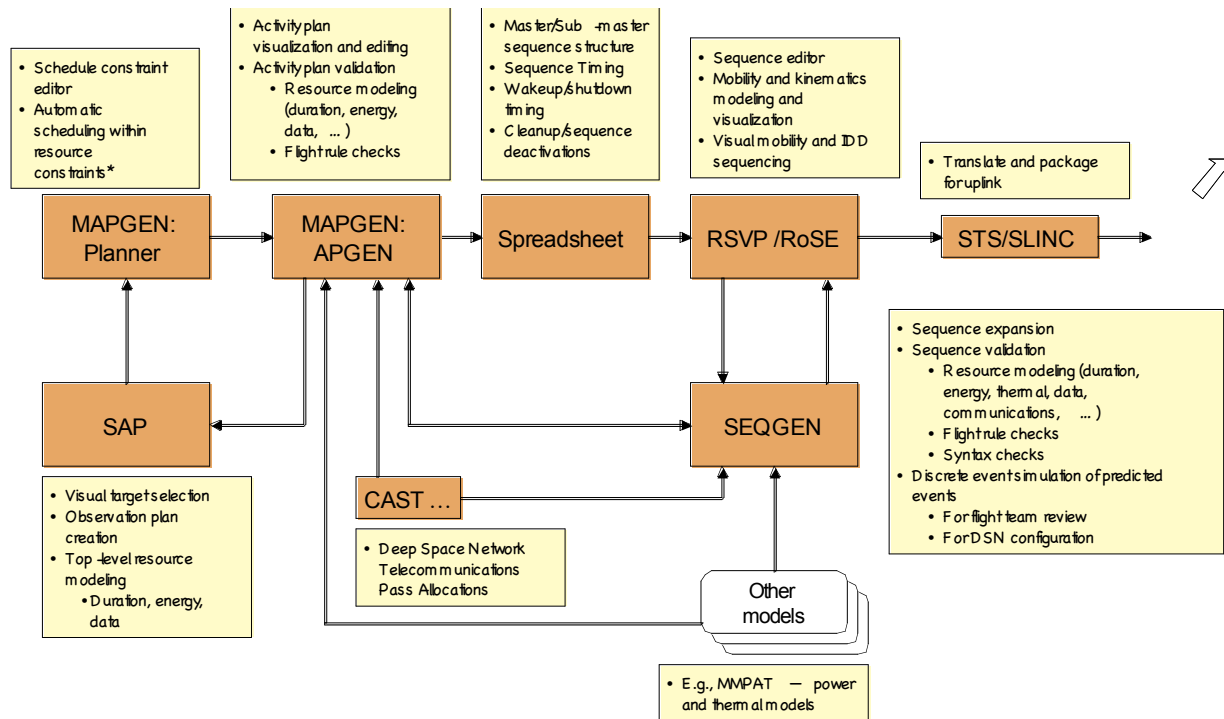


Figure 4. A Simplified View of the MER Uplink Process.

The MER uplink process (Fig. 4) is a sequential flow of products from tool to tool, usually with a Perl script between each pair, which translates products from one tool's format to the next. The process is repeated on a daily basis and begins with the downlink team having already processed the latest data and determined the current state of the spacecraft. Over the course of two years on Mars, the process has evolved significantly. Described here is the process as it exists today³.

B. Science Planning

The first order of business is science planning and the specification of science activities in the Science Activity Planner (SAP). Most significantly, SAP models the durations of activities and the volume of data each will produce. Within SAP, related activities are grouped together in Observations.

C. Activity Planning

Activity planning begins with a Perl script which translates from the SAP format to a format used by the MAPGEN activity planning and scheduling tool. The relationship between activities and observations, as specified in SAP, carries over into MAPGEN. In addition to modeling activity durations and data volumes as in SAP, MAPGEN models a number of power related resources such as battery state of charge, solar array current, and battery minimum usable amp-hours. Also modeled in MAPGEN are conflicts between activities which claim the same discrete resources. For example, Miniature Thermal Emission Spectrometer (Mini-TES) and Panoramic Camera (Pancam) activities both require the use of the Pancam Mast Assembly (PMA).

D. Sequencing

Sequencing begins with annotation of the sequence summary. The MAPGEN tool creates a time ordered list of activities, which a Perl script then translates into a file of tab-separated values called the sequence summary. The script assumes a one-to-one mapping of activities to sequences. Annotation of the sequence summary adds directives to the spreadsheet which specify absolute timing constraints. These annotations inform the subsequent translation from TSV to Rover Mark-Up Language (RML). RML is the XML file format used to specify sequences and commands. The translation from the sequence summary to RML is performed by another Perl script.

At this point the RML is ready to be refined using the Rover Sequence Editor (RSVP/RoSE). RSVP/RoSE allows users to instantiate commands from the command dictionary, piece them together in an ordered list, and place time constraints between adjacent commands.

Once complete, sequences in RML are translated, by a Perl script that invokes other software, into the actual uplink products radiated to the spacecraft.

IV. State Analysis Applied to MER

A. Opportunity Sol 63

In order to provide the richest set of MER activities from which to choose examples for this study, Opportunity sol 63 was chosen as a representative sol due to the wealth of diverse activities executed in that particular plan. The plan included mobility, use of the Instrument Deployment Device (IDD), a UHF communications relay with the Mars Odyssey Orbiter (ODY), pancam imaging, use of the Moessbauer Spectrometer (MB), remote sensing with the Mini-TES, CPU shutdown and wake-up, and direct-from-Earth communications using the high-gain antenna (HGA).

First we note some of the more visible state variables associated with a rover mission (Fig. 5), by starting with a position and heading state variable and applying the state discovery process described above. We then show other State Effects Diagrams (SED's) concerning some important subsystems, such as power (Fig. 11), mobility (Fig. 12), and data management (Fig. 8). Note the mixture of resource state variables such as power with activity state variables such as position and heading. We also note a common pattern for science instruments in Fig. 6.

Now that we have several state effects diagrams in place, we can focus our attention on a science activity and derive its elaboration from the SED. We will apply our general SED for a science instrument to the Microscopic Imager (MI). The MI is particularly interesting because it involves several diverse activities, some of which are quite analogous to other rover activities. The MI is a camera much like the pancam, navcams, and hazcams, which perform mosaics in much the same way. Furthermore, it requires the IDD to place the MI on the appropriate target.

We begin by applying the elaboration rules mentioned earlier to the SED for science instruments which applies to pancam and MI alike. The result is individual goals for MI stacks interleaved with goals on IDD position as shown in Fig. 9. Note the introduction of the concept of a macro-goal. A macro-goal is a goal used to group related goals, but does not apply to a specific state variable. We then apply more iterations of our elaboration rules to the MI stack sub-goals, obtaining the elaboration diagrams shown in Fig. 10 and Fig. 7.

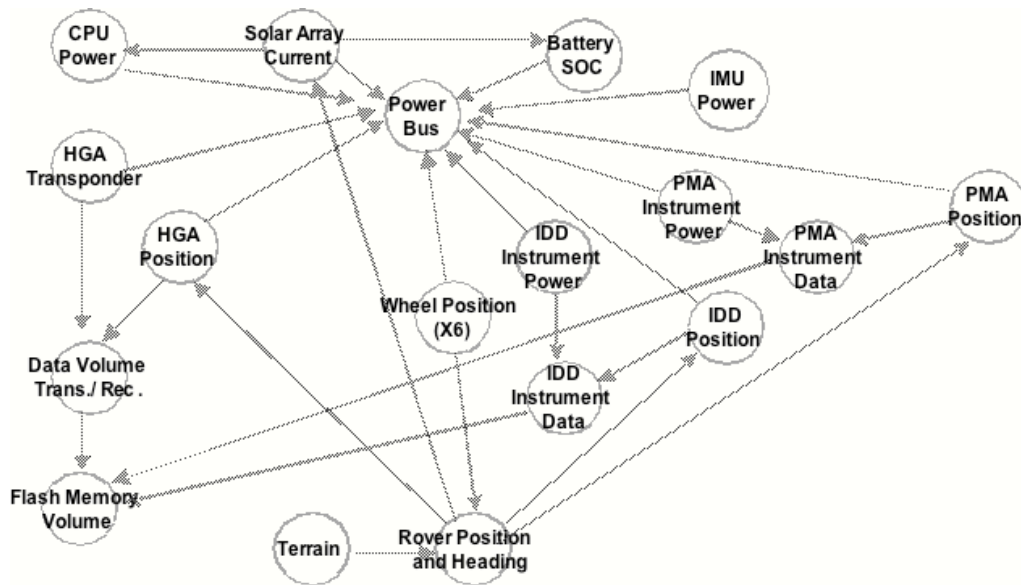


Figure 5. State Effects Diagram for Some Important MER State Variables.

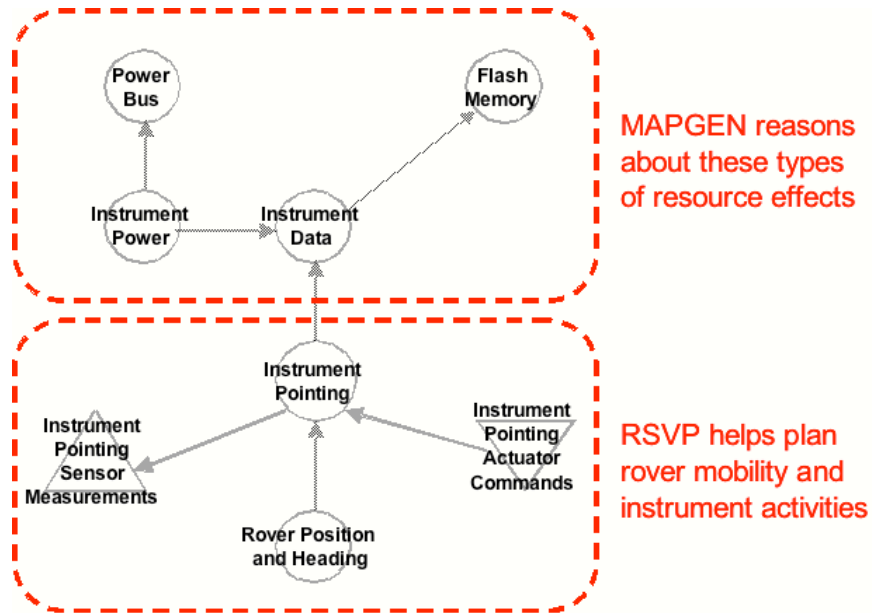


Figure 6. Typical State Effects Diagram for State Variables related to a Science Instrument.

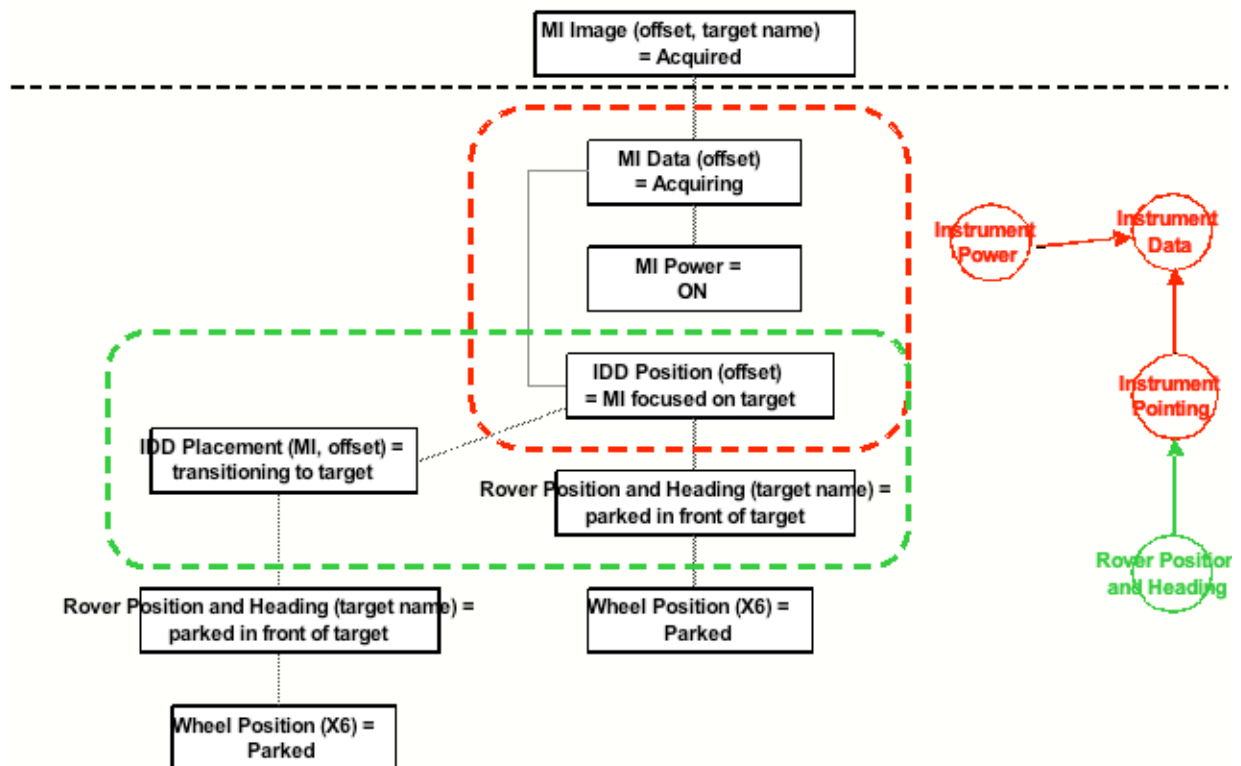


Figure 7. Elaboration Diagram for an MI Image Goal.

B. Goal-Based Operations

Goals represent activities, flight rules, and resources, therefore, actions and effects can be reasoned about side-by-side. Goals can overlap, because they are eventually merged by the planning system⁴. This allows for optimization by consolidation of redundant activities. State effects and elaborations allow the operators to trace resource violations back to the source. The mapping from activities to sequences is unclear. Sequence design patterns are not enforced much less standardized. This means that the state effects documented with state analysis are not enforced. Elaboration relationships allow for reusable patterns and enforce state effects, and macro-goals capture hierarchy design. Furthermore, the plan can be checked against an actual design document thus creating a direct verifiable link between design and operations.

Goal-based operations facilitates consistency across tools, eliminating the need for conversion scripts. In the current product-flow paradigm, products progress from one tool to the next through the exchange of data files. Reversing the flow through this process is awkward at best, and is usually avoided. Instead, fixes are often made in place without the benefit of the functionality of preceding tools.

Goals facilitate a shift to a work-flow paradigm (Fig. 13), where one product set managed by a common tool undergoes successive stages of refinement. Reversing to address problems is straightforward. Suppose we need to re-plan after sequences have been delivered, under the current approach used by MER, we would have to add an activity in SAP, plan and re-model resources in MAPGEN, perform the human translation from activities to sequences, insert an instantiation of the sequence into the sequence hierarchy, and create and edit the sequence in RSVP/RoSE. In a goal-based work-flow approach, we would simply add a goal, and let the tool elaborate the goal and model resources.

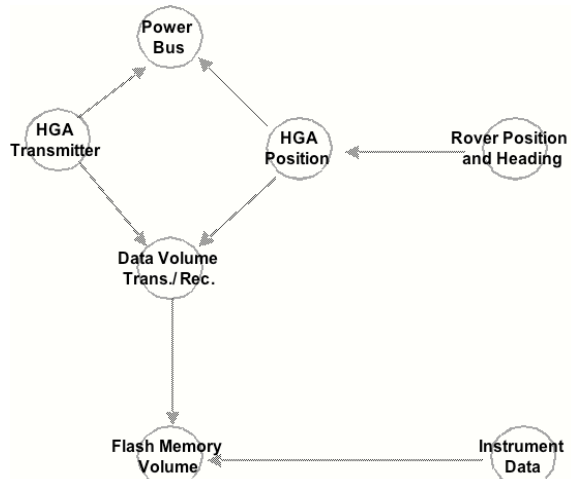


Figure 8. State Effects Diagram for Data-Related State Variables.

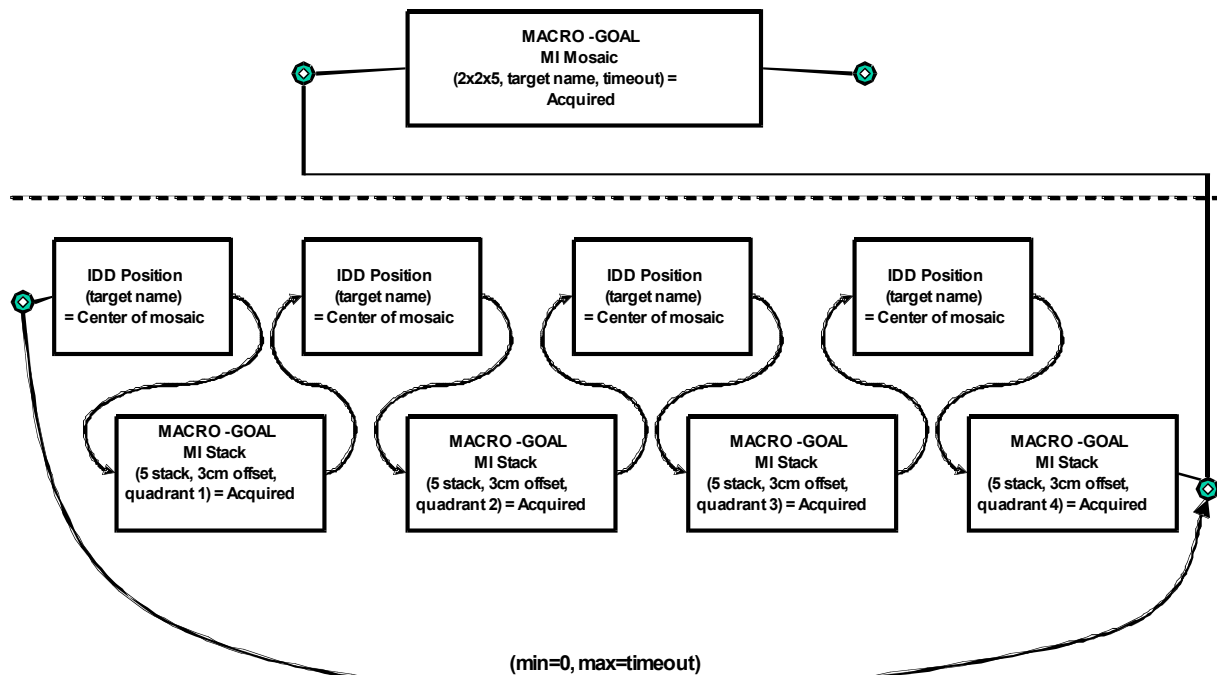


Figure 9. Elaboration Diagram for an MI Mosaic Goal.

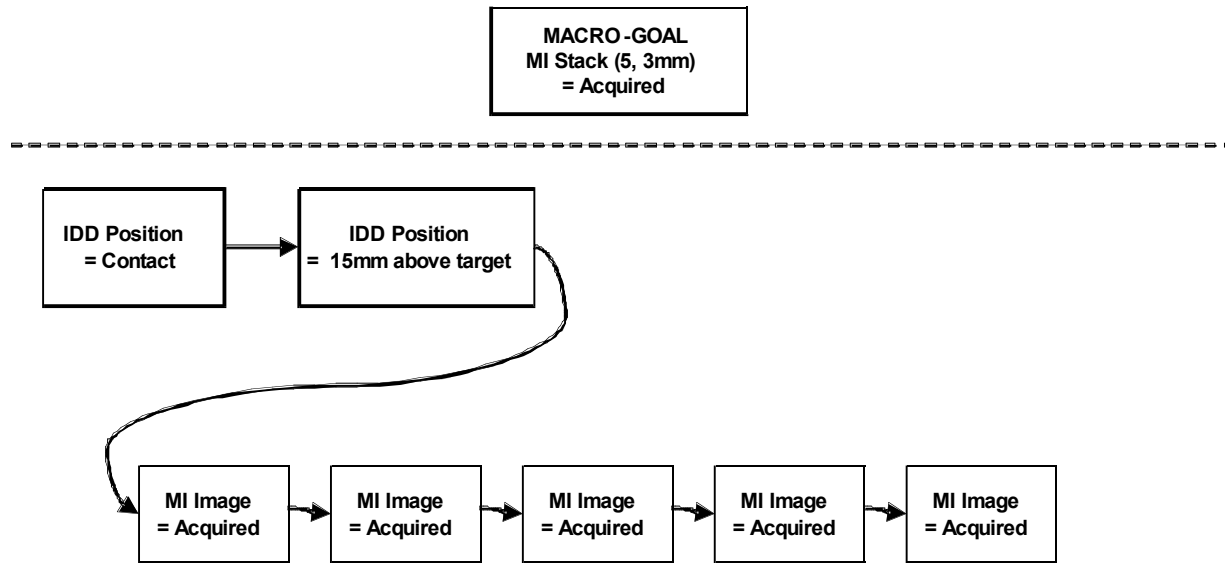


Figure 10. Elaboration Diagram for an MI Stack Goal.

V. Extreme MER

Some potential benefits of state analysis and goal-based operations go beyond the ambitions of the MER process. Theoretically, we can cut the MER uplink process in half. With goal-based operations, modeling and control are expressed in the same language (goals), eliminating the need for the translation from activities to sequences. These two pieces typically take up equal amounts of time within the process.

If a state-based software framework like MDS is used onboard the spacecraft, both the flight software (state variables) and the plan (goals) are built directly from the same model (state effects diagrams), enabling operators to more easily understand the flight software and what it will do with the commands they send, mitigating operator need for resident flight software experts, “hand-me-down” understanding, and/or guess work. Resource modeling (projection) can be done onboard, even during plan execution, so problems can be forestalled, as opposed to “I hope nothing happens on Mars that I didn’t model in MAPGEN”. This can prevent the loss of planning cycles used to recover from faults.

Can we use Projection for long-term strategic planning? Suppose we want to be at a particular outcrop in time for the weekend, such that we can build a plan which will cover several days, thus allowing the engineers to take time off. Suppose we also want to have enough onboard memory available for this lengthy science campaign. Or even suppose we want to have a particular image mosaic on the ground in time for a big press conference. Why not use MAPGEN to plan such things a week at a time? We don’t typically model these things so far in advance because the fidelity of our models, such as power, break down over time. But with state projection executing onboard the spacecraft, battery energy and solar array input measurements can inform planning onboard.

VI. Conclusion

The MER uplink process represents significant progress in operations design. The modeling of resources at the activity level, rather than the command level, allows operators to reason about the system at a higher level. As space mission systems become more complex, the number of states that an operator is forced to consider becomes prohibitive. Goal-based operations allows operators to focus at different levels of detail while explicitly considering complexity in terms of system state. State Analysis provides a systems engineering methodology and architecture which addresses Goal-based operations as an integrated part of system design.

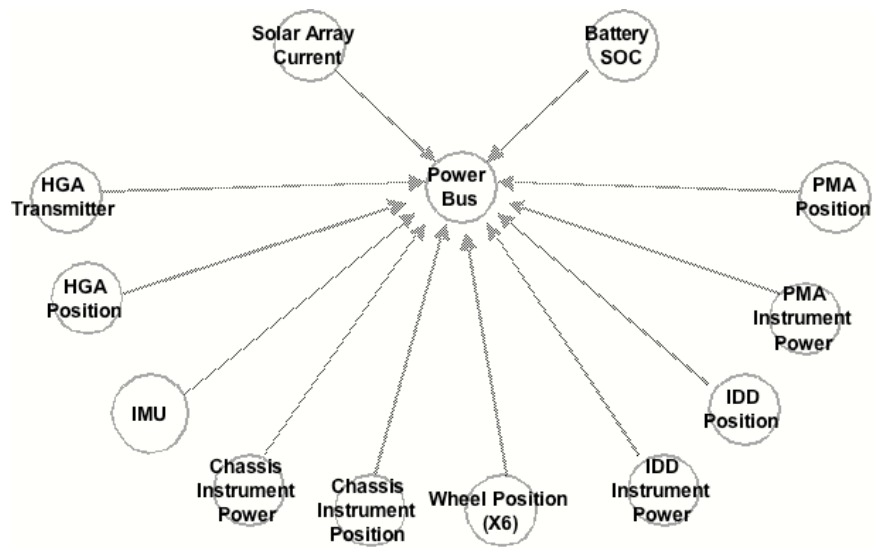


Figure 11. State Effects Diagram for Power-Related State Variables.

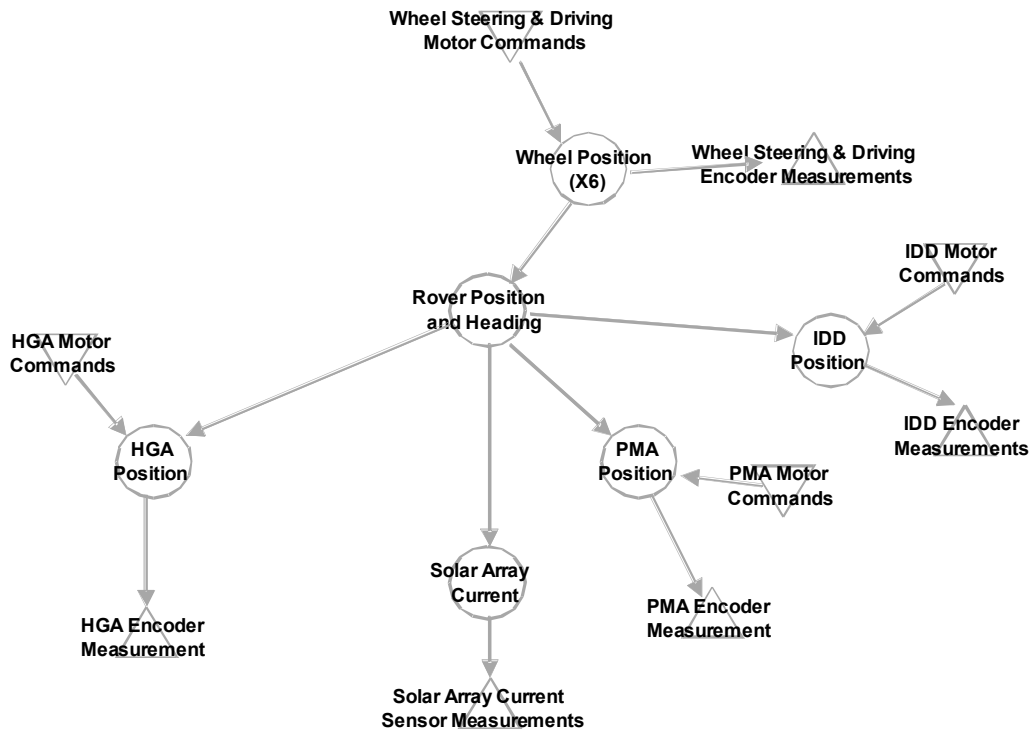


Figure 12. State Effects Diagram for Mobility-Related State Variables.

References

- ¹Ingham, M., Rasmussen, R., Bennett, M., and Moncada, A., "Engineering Complex Embedded Systems with State Analysis and the Mission Data System", *AIAA Journal of Aerospace Computing, Information and Communication*, Vol. 2, No. 12, Dec. 2005, pp. 507-536.
- ²Rasmussen, R., Ingham, M., Dvorak, D., "Achieving Control and Interoperability through Unified Model-based Systems and Software Engineering", Proceedings of Infotech 2005.
- ³Mishkin, A., Limonadi, D., Laubach, S., Bass, D., "Working the Martian Night Shift, the MER Surface Operations Process," *IEEE Robotics and Automation Magazine*, Vol. 13, No. 2, June 2006, pp. 46-53.
- ⁴Barrett, A., Knight, R., Morris, R., and Rasmussen, R., "Mission Planning and Execution Within the Mission Data System," Proceedings of the International Workshop on Planning and Scheduling for Space, 2004.
- ⁵Dvorak, D., Rasmussen, R., Reeves, G., and Sacks, A., "Software Architecture Themes in JPL's Mission Data System," Proceedings of the AIAA Guidance, Navigation, and Control Conference, number AIAA-99-4553, 1999.
- ⁶Dvorak, D., Rasmussen, R., and Starbird, T., "State Knowledge Representation in the Mission Data System," Proceedings of the IEEE Aerospace Conference, 2002.
- ⁷Rasmussen, R., "Goal-based Fault Tolerance for Space Systems Using the Mission Data System," Proceedings of the IEEE Aerospace Conference, 2001.

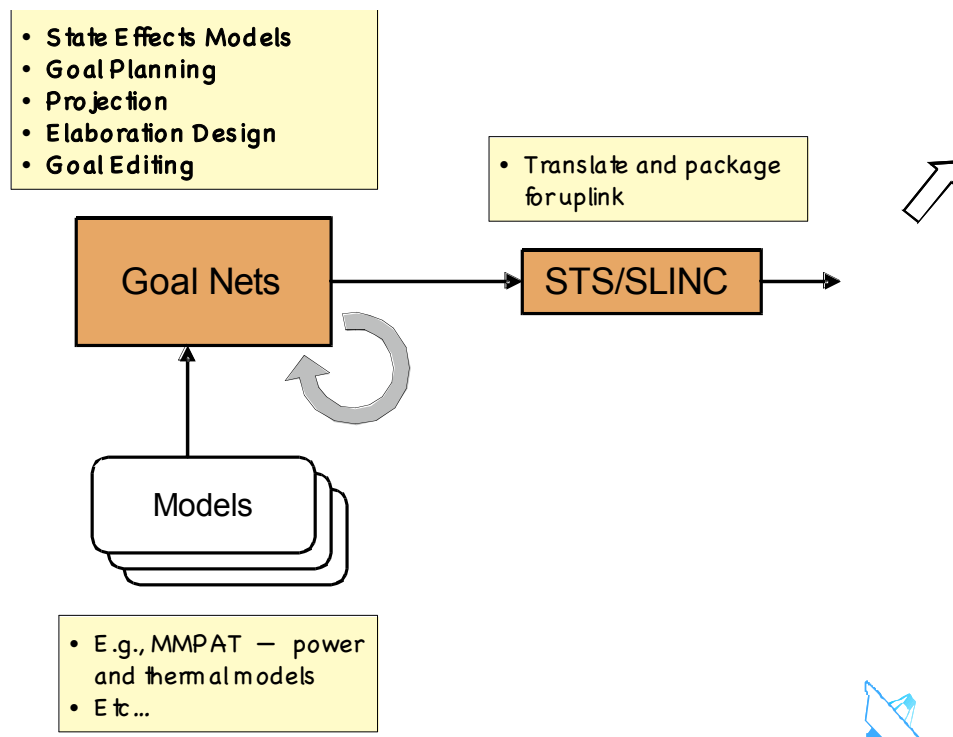


Figure 13. A Goal-Based Uplink Process.

Acknowledgments

The work described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.